

Workgroup: SPRING Working Group
Internet-Draft:
draft-geng-spring-sr-redundancy-protection-05
Published: 2 August 2021
Intended Status: Standards Track
Expires: 3 February 2022
Authors: X. Geng
Huawei Technologies
F. Yang, Ed.
Huawei Technologies
G. Mishra
Verizon Inc.

M. Chen
Huawei Technologies
P. Camarillo
Cisco Systems, Inc.

SRV6 for Redundancy Protection

Abstract

Redundancy Protection is a generalized protection mechanism to achieve the high reliability of service transmission in Segment Routing network. The mechanism inherits the "Live-Live" methodology, targeting to enhance the functionalities of Segment Routing over IPv6. Inspired by DetNet Packet Replication and Packet Elimination functions, two new Segments are introduced to provide replication and elimination functions on specific network nodes by leveraging SRV6 Segment programming capabilities.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in .

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 3 February 2022.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

- [1. Introduction](#)
- [2. Terminology](#)
 - [2.1. Requirements Language](#)
 - [2.2. Terminology and Conventions](#)
- [3. Redundancy Protection in Segment Routing Scenario](#)
- [4. Segment to Support Redundancy Protection](#)
 - [4.1. Redundancy Segment](#)
 - [4.2. Merging Segment](#)
- [5. Meta Data to Support Redundancy Protection](#)
- [6. Segment Routing Policy to Support Redundancy Protection](#)
- [7. IANA Considerations](#)
- [8. Security Considerations](#)
- [9. Acknowledgements](#)
- [10. References](#)
 - [10.1. Normative References](#)
 - [10.2. Informative References](#)
- [Authors' Addresses](#)

1. Introduction

Redundancy Protection is a generalized protection mechanism to achieve the high reliability of service transmission in Segment Routing network. Specifically, packets of flows are replicated at a network node into two or more copies, which are transported via different paths in parallel. When copies of packets are received and merged at one network node, the redundant packets are determined and further eliminated to guarantee only one copy of packets is transmitted. The mechanism inherits the "Live-Live" methodology, targeting to enhance the functionalities of Segment Routing over IPv6 [RFC8986]. Inspired by DetNet [RFC8655] Packet Replication and Packet Elimination Functions, two new Segments are introduced to provide the replication and elimination functions on specific

network nodes by leveraging SRv6 Segment programming capabilities. As it is unnecessary to perform switchover between different paths triggered by failure detection, redundancy protection can facilitate to achieve zero packet loss target when failure on either path happens.

Redundancy protection provides ultra reliable protection to many services, for example Cloud VR/Game, IPTV service and other type of video services, high value private line service etc. In this document, redundancy protection is applied to point-to-point service. The mechanism for P2MP service stays out of the scope of this document.

Segment Routing (SR) leverages the source routing paradigm. An ingress node steers a packet through an ordered list of instructions, called "segments". A segment can be associated to an arbitrary processing of the packet in the node identified by the segment.

This document extends the Segment Routing capabilities to support the redundancy protection in an SRv6 environment, including the definitions of two new Segments, meta data encapsulation, and a variation of Segment Routing Policy.

2. Terminology

2.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

2.2. Terminology and Conventions

SR: Segment Routing

URLLC: Ultra-Reliable Low-Latency Communication

VR: Virtual Reality

Red Node: Redundancy Node

Mer Node: Merging Node

FID: Flow IDentification

SN: Sequence Number

3. Redundancy Protection in Segment Routing Scenario

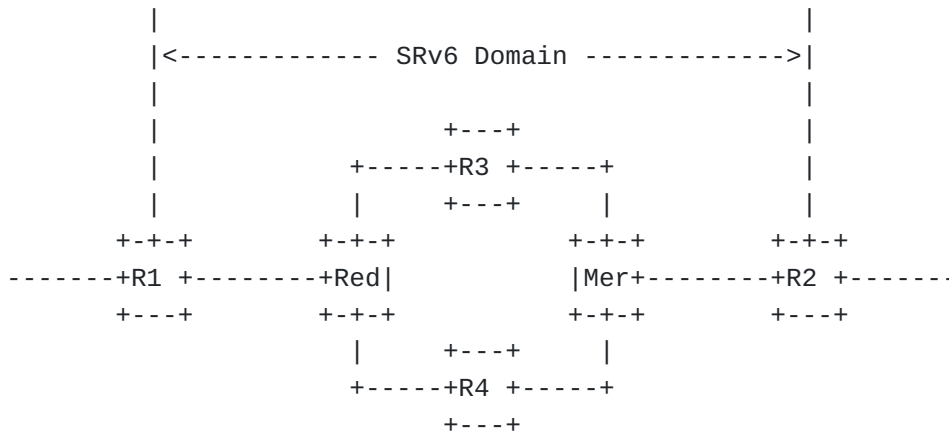


Figure 1: Example Scenario of Redundancy Protection in SRv6 Domain

This figure shows an example of redundancy protection used in SRv6 domain. R1, R2, R3, R4, Red and Mer are SR-capable nodes. When a flow is sent into SRv6 domain, the process is:

- 1) R1 receives the traffic flow and encapsulates packets with a list of segments destined to R2, which is instantiated as an ordered list of SRv6 SIDs.
- 2) When the packet flow arrives at Red node, known as Redundancy Node, each packet is replicated into two or more copies. Each copy of the packet is encapsulated with a new segment list, which represents different disjoint forwarding paths.
- 3) Meta data information such as flow identification (FID) and sequence number (SN) is used to facilitate the packet elimination on Merging node (Mer). Flow identification identifies the specific flow, and sequence number distinguishes the packet sequence of a flow. Meta data is either carried in the packet before it arrives at Red node, or added to each of the replicas at Red node.
- 4) The multiple replicas go through different paths until the Mer node. The first received packet of the flow is transmitted from Merging Node to R2, and the redundant packets are eliminated.
- 5) When there is any failures or packet loss in one path, the service continues uninterrupted through the other path without break.
- 6) Sometimes, the packet will arrive out of order because of redundancy protection, the function of reordering may be also necessary on Merging Node. In such case the Merging node may include a reordering function, which is implementation specific and out of the scope of this document.

In this example, service protection is supported by utilizing two packet flows transmitted over two forwarding paths. It is noted that there is no limitation of the number of replicas. For a unidirectional flow, Red node supports replication function, and Mer node supports elimination function. Reordering function MAY be required in combination of elimination function on merging node. To minimize the jitter caused by random packet loss, the disjoint paths are recommended to have similar path forwarding delay.

4. Segment to Support Redundancy Protection

To achieve the packet replication and elimination functions, Redundancy Segment and Merging Segment, as well as the related SRv6 Endpoint Behavior are introduced.

4.1. Redundancy Segment

Redundancy Segment is the identifier of packets which need the replication function on redundancy node. It is also a variation of Binding SID, and associated with a Redundancy Policy to provide segment lists of disjoint paths. Thus, Redundancy segment is associated with service instructions, indicating the following operations:

- *Steers the packet into the corresponding redundancy policy

- *Encapsulates flow identification and sequence number in packets if the two information is not carried in packets

- *Packet replication and segment encapsulation based on the information of redundancy policy, e.g., the number of replication copies, an ordered list of segments with a topological instruction

In the case of SRv6, a new behavior End.R for Redundancy Segment is defined. An instance of a redundancy SID is associated with a redundancy policy B and a source address A. In the following description, End.R behavior is specified in the encapsulation mode. The End.R behavior in the insertion mode is for further study.

When an SRv6-capable node (N) receives an IPv6 packet whose destination address matches a local IPv6 address instantiated as an SRv6 SID (S), and S is a Redundancy SID, N does:

```

S01. When an SRH is processed {
S02.   If (Segments Left>0)   {
S03.     Decrement IPv6 Hop Limit by 1
S04.     Decrement Segments Left by 1
S05.     Update IPv6 DA with Segment List[Segments Left]
S06.     Add flow identification and sequence number if indicated*
S07.     Duplicate the packets (as number of active SID lists in B)
S08.     Push the new IPv6 headers to each replica. The IPv6 header
           contains an SRH with the SID list in B
S09.     Set the outer IPv6 SA to A
S10.     Set the outer IPv6 DA to the first SID of new SRH SL
S11.     Set the outer Payload Length, Traffic Class, Flow Label,
           Hop Limit and Next-Header fields
S12.     Submit the packet to the egress IPv6 FIB lookup
           for transmission to the new destination
S13.   }
S14. }

```

* Adding flow identification and sequence number is an optional behavior for Redundancy Segment. The instruction execution is determined and explicitly indicated by SR policy or Segment itself.

4.2. Merging Segment

Merging Segment is associated with service instructions, indicates the following operations:

*Packet merging and elimination: forward the first received packets and eliminate the redundant packets

In order to eliminate the redundant packet of a flow, merging node utilizes sequence number to evaluate the redundant status of a packet. Note that implementation specific mechanism could be applied to control the amount of state monitored on sequence number, so that system memory usage can be limited at a reasonable level.

As merging node needs to maintain the state of flows, a centralized controller should have a knowledge of merging nodes capability, and never provision the redundancy policy to redundancy node when the computation result goes beyond the flow recovery capability of merging node. The capability advertisement of merging node will be specified separately elsewhere, which is not within the scope of this document.

In the case of SRv6, a new behavior End.M for Merging Segment is defined.

When an SRv6-capable node (N) receives an IPv6 packet whose destination address matches a local IPv6 address instantiated as an SRv6 SID (S), and S is a Merging SID, N does:

```

S01. When an SRH is processed {
S02.   If (Segments Left > or == 0) {
S03.     Acquire the sequence number of received packet and
        look it up in table
S04.       If (this sequence number does not exist in the table) {
S05.         Store this sequence number in table
S06.         Remove the outer IPv6+SRH header
S07.         Decrement IPv6 Hop Limit by 1 in inner SRH
S08.         Decrement Segments Left by 1 in inner SRH
S09.         Update IPv6 DA with Segment List[Segments Left] in inner SRH
S10.         Submit the packet to the egress IPv6 FIB lookup and transmit
S11.       }
S12.     ELSE {
S13.       Drop the packet
S14.     }
S15.   }
S16. }

```

5. Meta Data to Support Redundancy Protection

To support the redundancy protection function, flow identification and sequence number are required. Flow identification identifies one specific flow of redundancy protection, and is usually allocated from centralized controller to the SR ingress node or redundancy node in SR network. Sequence number distinguishes the packets within a flow by specifying the order of packets. It is usually generated at SR ingress node. If necessary, redundancy node can also facilitate to add sequence number if required. Thus, encapsulations of flow identification and sequence number should be specified accordingly.

6. Segment Routing Policy to Support Redundancy Protection

Redundancy Policy is a variation of SR Policy to conduct the replicas to multiple disjoint paths for redundancy protection. It extends SR policy to include more than one ordered lists of segments between redundancy node and merging node, and all the ordered lists of segments are used at the same time to steer the copies of flow into different disjoint paths.

7. IANA Considerations

This document requires registration of End.R behavior and End.M behavior in "SRv6 Endpoint Behaviors" sub-registry of "Segment Routing Parameters" registry.

8. Security Considerations

TBD

9. Acknowledgements

The authors would like to thank Bruno Decraene, Ron Bonica, James Guichard, Jeffrey Zhang, Balazs Varga for their valuable comments and discussions.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.

10.2. Informative References

- [RFC8655] Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", RFC 8655, DOI 10.17487/RFC8655, October 2019, <<https://www.rfc-editor.org/info/rfc8655>>.

Authors' Addresses

Xuesong Geng
Huawei Technologies
China

Email: gengxuesong@huawei.com

Mach(Guoyi) Chen
Huawei Technologies
China

Email: mach.chen@huawei.com

Fan Yang
Huawei Technologies
China

Email: shirley.yangfan@huawei.com

Pablo Camarillo Garvia
Cisco Systems, Inc.
Spain

Email: pcamaril@cisco.com

Gyan Mishra
Verizon Inc.

Email: gyan.s.mishra@verizon.com